

AUG 29 2005

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Pamela A. Binns et al.

Title: METHODS AND APPARATUS FOR SHARING SLACK IN A TIME-PARTITIONED SYSTEM

AP
2127
JPW

Docket No.: H16-25537.1-1611 (256.048US1)
Filed: December 29, 2000
Examiner: Niles R Shah

Serial No.: 09/751,834
Due Date: September 1, 2005
Group Art Unit: 2127

MS Appeal Brief - Patents

Commissioner for Patents

P.O. Box 1450

Alexandria, VA 22313-1450


We are transmitting herewith the following attached items (as indicated with an "X"):

- ☒ Return postcard.
- ☒ Substitute Appeal Brief Under 37 C.F.R. 41.37 (29 pgs.).

If not provided for in a separate paper filed herewith, Please consider this a PETITION FOR EXTENSION OF TIME for sufficient number of months to enter these papers and please charge any additional fees or credit overpayment to Deposit Account No. 19-0743.

SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.

Customer Number 21186


By: 
Atty: Bradley A. Forrest
Reg. No. 30,837

CERTIFICATE UNDER 37 CFR 1.8: The undersigned hereby certifies that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail, in an envelope addressed to: MS Appeal Brief - Patents, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on this 25 day of August, 2005.

CANDIS BUENDING

Name

Signature



SCHWEGMAN, LUNDBERG, WOESSNER & KLUTH, P.A.

(GENERAL)

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Application of:)

Pamela A. Binns)

Examiner: Nilesh R. Shah

Serial No.: 09/751,834)

Group Art Unit: 2127

Filed: December 29, 2000)

Docket: H16-25537.1-1611

For: METHODS AND APPARATUS FOR SHARING SLACK IN A TIME
PARTITIONED SYSTEM)

SUBSTITUTE APPEAL BRIEF UNDER 37 C.F.R. § 41.37

Mail Stop Appeal Brief- Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

This Substitute Appeal Brief is presented in support of the Notice of Appeal to the Board of Patent Appeals and Interferences, filed on March 18, 2005, from the Final Rejection of claims 1-30 of the above-identified application, as set forth in the Final Office Action mailed on November 18, 2004.

This Substitute Appeal Brief is being filed in response to the Notification of Non-Compliant Appeal Brief mailed August 1, 2005, which noted the lack of Appendices for Related Proceedings and Evidence. These Appendices are attached hereto. No other amendments to the brief have been made herein.

The Appellants respectfully request consideration and reversal of the Examiner's rejections of pending claims.

It is believed that no further filing fee is required. However, any fees deemed to be due may be charged to Deposit Account No. 19-0743.



APPEAL BRIEF UNDER 37 C.F.R. § 41.37

TABLE OF CONTENTS

	<u>Page</u>
1. REAL PARTY IN INTEREST	2
2. RELATED APPEALS AND INTERFERENCES.....	3
3. STATUS OF THE CLAIMS	4
4. STATUS OF AMENDMENTS	5
5. SUMMARY OF THE INVENTION	6
6. ISSUES PRESENTED FOR REVIEW	8
7. ARGUMENT	9
8. SUMMARY	22
APPENDIX I-The Claims on Appeal	23
APPENDIX II - Related Proceedings	28
APPENDIX III - Evidence.....	29

Serial No.: 09/715,834

Filing Date: December 29, 2000

Title: METHODS AND APPARATUS FOR SHARING SLACK IN A TIME PARTITIONED SYSTEM

1. REAL PARTY IN INTEREST

The real party in interest of the above-captioned patent application is the assignee,
HONEYWELL INTERNATIONAL INC..

2. RELATED APPEALS AND INTERFERENCES

This application is related to Application Serial No.: 09/751,955, filed December 29, 2000, and titled: METHODS AND APPARATUS FOR SLACK STEALING WITH DYNAMIC THREADS, which is currently under appeal. There are no other appeals or interferences known to Appellant that will have a bearing on the Board's decision in the present appeal.

3. STATUS OF THE CLAIMS

Claims 1-30 are pending in the application and all are the subject of the present appeal.

Serial No.: 09/715,834

Filing Date: December 29, 2000

Title: METHODS AND APPARATUS FOR SHARING SLACK IN A TIME PARTITIONED SYSTEM

4. STATUS OF AMENDMENTS

No claims were amended after the final rejection.

5. SUMMARY OF THE INVENTION

As described in claim 1, a data processing system executes tasks in a different time partitions (103 in FIG. 18A, Paragraph [00156] et seq.) Available slack is determined (105) and allocated to tasks in different time partitions (111 in FIG. 18B, [00436] to [00443]). In claim 2, the tasks may be aperiodic non-essential tasks (111). Claim 3 indicates that the tasks comprise essential and non-essential tasks (103) and tasks that are allocated slack are new non-essential tasks and enhancements to essential tasks (111). In claim 4, both timeline slack and reclaimed slack are determined (105, [00103], [00438]).

Independent claim 5 is a machine readable medium (132) with elements similar to claim 1. It further recites that tasks are scheduled to execute in different time partitions (111, [00016], [00441]).

Independent claim 6 also describes a data processing system that executes tasks in different time partitions (103 in FIG. 18A, Paragraph [00156] et seq.) Unscheduled execution time is collected from at least one time partition (105) and is allocated to a task in a different time partition (111 in FIG. 18B, [00436] to [00443]). In claim 7, the tasks may be aperiodic non-essential tasks (111). Claim 8 indicates that the tasks comprise essential and non-essential tasks (103) and tasks that are allocated slack are new non-essential tasks and enhancements to essential tasks (111). In claim 9, both timeline slack and reclaimed slack are determined (105, [00438]).

Independent claim 10 is a machine-readable medium claim corresponding to claim 5.

Independent claim 11 recites a method in a time-partitioned system that executes essential and non-essential tasks (103). Available slack is determined from timeline slack and reclaimed slack (105), and is polled in a common slack pool (109). The slack is then allocated to tasks from the common slack pool (111). Claim 12 allocates slack to non-essential tasks (111). Claim 13 allocates slack to new non-essential tasks and enhancement to essential tasks (103).

Independent claim 14 has elements similar to claim 11, and further describes scheduling tasks to execute in different time partitions (103 in FIG. 18A, Paragraph [00156] et seq.).

Independent claim 15 is also similar to claim 11, and includes allocating slack to a task in any time partition (111). Claim 16 allocates slack to non-essential tasks (111). Claim 17 allocates slack to new non-essential tasks and enhancement to essential tasks (103).

Independent claim 18 is a machine-readable medium version of claim 15.

Independent claim 19 describes a time-partitioned system (103 in FIG. 18A, Paragraph [00156] et seq.) with a processor that executes a plurality of essential and non-essential tasks (103). Each of the tasks has associated with it at least one worst case execution time ([0017] Figure 7). An executive ([0081] 17) is in communication with the processor and controls dispatching of tasks on the processor. The executive ([0081] 17) has a first module [0017] that determines available slack and a second module [0017] that allocates the available slack to tasks in different time partitions (111 in FIG. 18B, [00436] to [00443]). In claim 20, the first module determines available slack by determining slack from the group consisting of timeline slack, reclaimed slack, and idle time (105, [00438]). Claim 21 adds that the first module maintains a pool of available slack (111). In claim 22 the first module maintains a common pool of available slack that can be used by tasks in any time partition partitions (103 in FIG. 18A, Paragraph [00156] et seq.). In claim 23, the second module allocates available slack to tasks that are non-essential (111). In claim 24, such tasks are new non-essential tasks and enhancements to essential tasks (111). In claim 25, the executive further has a third module that assigns different priority levels to tasks [0087]. In claim 26, the first module determines available slack for tasks at each priority level ([00107-00147], [309]). In claim 27, the second module allocates available slack to tasks in order of priority [00309]. In claim 28, the system is a flight control system (FIG. 1, [00073-00075]). In claim 29, the system is a real-time control system [00074]. In claim 30, the executive comprises a single set of slack variables [00120-00132] and a single slack table (FIG. 20E).

6. ISSUES PRESENTED FOR REVIEW

Claims 1, 5, 11 and 18 were rejected under the judicially created doctrine of obviousness-type double patenting over claim 1 of U.S. Patent Application No. 09/751,955.

Claims 1-30 were rejected under 35 USC § 103(a) as being unpatentable over Atlas et al. ('Slack Stealing Job Admission Control') in view of Ezer et al. (U.S. Patent No. 6,275,239).

7. ARGUMENT

1) The Applicable Law

To establish an obviousness-type double patenting rejection, the Examiner has the burden to show that (1) the inventions claimed (2) are not patentably distinct and (3) are based on a prima facie showing of obviousness. This analysis must be based on what the claim defines and not on the claim language itself, as required by the Federal Circuit:

[I]t is important to bear in mind that comparison can be made only with what invention is *claimed* in the earlier patent, paying careful attention to the rules of claim interpretation to determine what invention a claim *defines* and not looking to the claim language for anything that happens to be mentioned in it as though it were a prior art reference. ... [W]hat is claimed is what is *defined by the claim taken as a whole*, every claim limitation ... being material. *General Foods Corp. V. Studiengesellschaft Kohle mbH*, 972 F.2d 1272, 23 USPQ 2d, 1839, 1845 (Fed. Cir. 1992). (emphasis in original.)

The Examiner has the burden under 35 U.S.C. 103 to establish a *prima facie* case of obviousness. *In re Fine*, 837 F.2d 1071, 1074, 5 USPQ2d 1596, 1598 (Fed. Cir. 1988). As part of establishing a *prima facie* case of obviousness, the Examiner must show that some objective teaching in the prior art or some knowledge generally available to one of ordinary skill in the art would lead an individual to combine the relevant teaching of the references. *Id.*

The court in *Fine* stated that:

Obviousness is tested by "what the combined teaching of the references would have suggested to those of ordinary skill in the art." *In re Keller*, 642 F.2d 413, 425, 208 USPQ 871, 878 (CCPA 1981)). But it "cannot be established by combining the teachings of the prior art to produce the claimed invention, absent some teaching or suggestion supporting the combination." *ACS Hosp. Sys.*, 732 F.2d at 1577, 221 USPQ at 933. And "teachings of references can be combined *only* if there is some suggestion or incentive to do so."

Id. (emphasis in original).

The M.P.E.P. adopts this line of reasoning, stating that:

“In order for the Examiner to establish a *prima facie* case of obviousness, three base criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, and not based on Appellant's disclosure. *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed.Cir. 1991)”. *M.P.E.P.* 2142. *In re Royka*, 490 F.2d 981, 180 USPQ 580 (CCPA, 1974). *MPEP* §2143.

The Office Action must provide specific, objective evidence of record for a finding of a suggestion or motivation to combine reference teachings and must explain the reasoning by which the evidence is deemed to support such a finding. *In re Sang Su Lee*, 277 F.3d 1338, 61 USPQ2d 1430 (Fed. Cir. 2002).

The test for obviousness under § 103 must take into consideration the invention as a whole; that is, one must consider the particular problem solved by the combination of elements that define the invention. *Interconnect Planning Corp. v. Feil*, 774 F.2d 1132, 1143, 227 USPQ 543, 551 (Fed. Cir. 1985). The Examiner must, as one of the inquiries pertinent to any obviousness inquiry under 35 U.S.C. § 103, recognize and consider not only the similarities but also the critical differences between the claimed invention and the prior art. *In re Bond*, 910 F.2d 831, 834, 15 USPQ2d 1566, 1568 (Fed. Cir. 1990), *reh'g denied*, 1990 U.S. App. LEXIS 19971 (Fed. Cir. 1990). Finally, the Examiner must avoid hindsight. *Id.*

An invention can be obvious even though the suggestion to combine prior art teachings is not found in a specific reference. *In re Oetiker*, 24 USPQ2d 1443 (Fed. Cir. 1992). At the same time, however, although it is not necessary that the cited references or prior art specifically suggest making the combination, there must be some teaching somewhere which provides the suggestion or motivation to combine prior art teachings and applies that combination to solve the same or similar problem which the claimed invention addresses. One of ordinary skill in the art will be presumed to know of any

such teaching. (See, e.g., *In re Nilssen*, 851 F.2d 1401, 1403, 7 USPQ2d 1500, 1502 (Fed. Cir. 1988) and *In re Wood*, 599 F.2d 1032, 1037, 202 USPQ 171, 174 (CCPA 1979)).

If an independent claim is nonobvious under 35 U.S.C. §103, then any claim depending therefrom is nonobvious. MPEP §2143.03.

2) *Discussion of the Rejection of Claims Under the Judicially Created Doctrine of Obviousness Type Double Patenting.*

Claims 1, 5, 11 and 18

Claims 1, 5, 11 and 18 were rejected under the judicially created doctrine of obviousness type double patenting over claim 1 of U.S. Patent Application No. 09/751,955. This rejection is respectfully traversed, because the Office Action specifically indicates that the elements of the claims are different in “the claimed way the slack is allocated.” Thus, a prima facie showing of obviousness has not been established, and the rejection should be withdrawn. Further, the rejection is not ripe for examination, as U.S. Patent Application No. 09/751,955 has not yet issued. Thus, the claims are not being rejected in view of a claim in an issued patent.

3) *Discussion of the Rejection of the Claims Under 35 U.S.C. § 103(a) as being unpatentable over Atlas et al. (‘Slack Stealing Job Admission Control’) in view of Ezer et al. (U.S. Patent No. 6,275,239).*

Claims 1-30

Claims 1-30 were rejected under 35 USC § 103(a) as being unpatentable over Atlas et al. (‘Slack Stealing Job Admission Control’) in view of Ezer et al. (U.S. Patent No. 6,275,239). This rejection is respectfully traversed, as the references either alone or combined do not teach or suggest each and every element of the claimed invention.

The rejection is respectfully traversed on the basis that the combined references do not teach or suggest each and every element, and that the references are not properly

combinable for at least one of several reasons including, they are directed at different problems, there is no proper suggestion to combine them, the combination is based on hindsight, and a reasonable expectation of success in making the combination has not been established.

Atlas discloses a slack-stealing job admission control system for scheduling periodic firm-deadline tasks with variable resource requirements. See Abstract. The Examiner admits that, and Applicants are unable to find mention of any teaching of the use of different time partitions.

Ezer discloses a media coprocessor for performing 3D graphics, video, and audio functions. See Abstract. An update interval, synchronized to a video frame, is divided into a number of partitions for sequential processing of audio, video, and 3D graphics. Applicants are unable to find any mention whatsoever of “slack” or “allocating slack” within Ezer.

Regarding independent claim 1, the Examiner concedes that Atlas fails to disclose the use of setting different time partitions. However, the Examiner states it would have been obvious to one skilled in the art at the time of the invention to combine the teachings of Atlas and Ezer to ensure that different time partitions have access to slack. The Examiner argues that by being able to allocate slack to different time partitions, the user can determine which tasks get the slack first, thus making the entire system more efficient.

The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, not in applicant’s disclosure. *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991); MPEP § 2143. The argument presented by the Examiner to combine the references is clearly based on the claimed combination. In other words, the Examiner takes the claimed combination, and it’s attendant benefits, and then picks features from unrelated references, combines them, and states that it provides the benefits of the claimed invention. Nowhere is a suggestion to make the combination found in the prior art. Rather, the benefit provided by the combination itself is used as the suggestion. This is clearly improper, and the rejection should be withdrawn.

In the Response to Arguments section of the Final Office Action, the Examiner indicates that “Ezer also provides motivation as to why one would want to include the use of time-partitioned tasks in an invention (col. 1, line 60 – col. 2, line 7).” This language is reproduced here for convenience:

“In contrast, the present invention provides a much less expensive computing system without sacrificing much in the way of functionality, quality, and versatility. The present invention achieves this by designing a single, integrated media co-processor chip which is capable of processing graphics, video and audio. This reduces costs because it minimizes duplicate functionalities and redundant circuitry. With the media co-processor chip, the present invention partitions different media tasks so that they are performed sequentially in a time-division multiplexed format. For a given amount of time synchronized to the video frame, a specified amount of time is partitioned to perform audio tasks; a specified amount of time is partitioned to process video; and the time left over is partitioned to render graphics.”

Applicants fail to see how this language suggests the inclusion of time partitions in a slack stealing algorithm. Rather, it appears to point to dividing media tasks into partitions, and then timing execution of the partitions so that all are completed for a video frame. There is no indication that additional tasks need to be scheduled or considered, and no teaching to suggest that it would be applicable in a slack stealing algorithm, since the times for each partition appear to be fixed. It appears to be directed to an entirely different problem, which further points away from finding a suggestion to combine it with Atlas.

Independent claim 1 recites *inter alia* a method of allocating slack to tasks in different time partitions in a data processing system executing tasks in different time partitions.

Applicants respectfully assert that the Examiner fails to provide any specific factual basis within either Atlas or Ezer that would provide any credible motivation to combine these two complex inventions in the manner suggested by the Examiner.

Moreover, Applicants respectfully points out that the **Examiner has failed to provide any specific basis (i.e. column or page, and line number) within either reference for the limitation “allocating slack to tasks in different time partitions”**. Thus, even if the references are combined, the combination does not teach or suggest each and every claim element. To establish a *prima facie* case of obviousness under 35 U.S.C. §103, the prior art reference (or references when combined) must teach or suggest every limitation of the claim. *In re Royka*, 490 F.2d 981, 180 USPQ 580 (CCPA, 1974). MPEP §2143.

The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, and not based on applicant's disclosure. *M.P.E.P.* § 2142 (citing *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed.Cir. 1991)). **No reasonable expectation of success in combining the references has been pointed to by the Examiner.** Ezer appears to contain no disclosure whatsoever concerning “slack” or “allocating slack”. Nor, as the Examiner concedes, does Atlas appear to contain any disclosure concerning “executing tasks in different time partitions”. Thus, there is no indication within the references themselves that combining them would succeed in producing the claimed invention.

These references describe fairly complex algorithms for very different computer science concepts. Ezer describes different partitions for processing very different types of data, including audio data, video data and 3-D graphics on a single media coprocessor chip. The partitions are synchronized to process each type of data for and during a single video frame. This is a very narrow graphics oriented computer processing concept. Atlas describes a complex algorithm for slack stealing involving a slack stealing job admission control, which is a generalization of Rate Monotonic Scheduling and a semiperiodic task model, involving probability density functions and independent and identically distributed random variables for resource requirements for jobs. See pages 2-4. It is a radically different approach to a very different problem from Ezer. The Examiner has not pointed out where the references provide a reasonable expectation of success in combining them, much less an actual suggestion to combine them based on the teaching of the prior art. Thus, a *prima facie* case of obviousness has not been established, and the rejection should be withdrawn.

Ezer processes audio, video and graphics in three different time partitions for each separate video frame. Claim 1 recites allocating slack to tasks in different time partitions. The partitions in Ezer appear to be timed to allow complete processing for each video frame, thus obviating any need for calculating or sharing slack. Further, it would appear to be difficult to utilize slack from different time partitions in Ezer, since the hardware used for each partition is different.

Therefore, Applicants respectfully assert that a *prima facie* case of obviousness has not been established. First, the Examiner has not provided any teaching, suggestion, or motivation to combine the references from the prior art. Second, the Examiner has not provided any credible teaching, suggestion, or motivation in the knowledge generally available to one of ordinary skill in the art, to combine the Atlas and Ezer references to arrive at Applicants' claimed subject matter. Third, each and every element has not been shown in the references, alone or combined, and fourth, no reasonable expectation of success in combining the references has been established.

For the above reasons, independent claim 1 should be found to be allowable over any combination of Atlas and Ezer, and Applicants respectfully request that the rejection of claim 1 under 35 U.S.C. §103(a) as being unpatentable over Atlas in view of Ezer should be reversed.

If an independent claim is nonobvious under 35 U.S.C. §103, then any claim depending therefrom is nonobvious. MPEP §2143.03.

Claims 2-4, which depend from claim 1, directly or indirectly, and incorporate all of the limitations therein, are also asserted to be allowable for the reasons presented above.

Further, claim 2 recites that the tasks that are allocated slack are aperiodic, non-essential tasks. Atlas appears to be limited to only a static set of execution threads, i.e. a fixed set of recurring tasks without any new periodic tasks being activated and without any periodic tasks being deactivated. However, claim 2 recites that slack is allocated to aperiodic, non-essential tasks. The Final Office Action asserts that this is taught in Atlas at page 4 lines 19, 23-24, and 36-37. Atlas, references aperiodic tasks in passing comments about other work, but specifically indicates that "We use slack stealing to

admit jobs of variable length.” And “For SSJAC, we do not use any aperiodic servers,” in section 3.2, paragraphs 1 and 3 on page 4. Thus, it does not teach the limitations of claim 2 and the rejection should be reversed.

Claim 3 recites that tasks are allocated slack are from the group consisting of new non-essential tasks and enhancements to essential tasks. The Final Office Action indicates that this is shown in Atlas at page 4 lines 19, 23-24, and 36-37. No such teaching is seen at the referenced lines. Only periodic jobs of variable length appears to be referenced, as well as various priorities of jobs. Since this element appears to be lacking from the references, the rejection should be reversed.

Claim 4 recites that both timeline slack and reclaimed slack are determined. The Final Office Action references page 4 lines 19, 23-24, and 36-37, and page 5, lines 13-17. These sections of Atlas have been reviewed, and no reference to those types of slack were found.

Independent claims 5-6, 10-11, 14-15, 18, and 19, along with their associated dependent claims, were rejected essentially based upon the same grounds as independent claim 1. Applicants respectfully assert that these claims are all patentable over any suggested combination of Atlas and Ezer for the reasons presented earlier regarding independent claim 1. Applicants respectfully requests that the rejection of claims 1-30 under 35 U.S.C. §103(a) as being unpatentable over Atlas in view of Ezer should be reversed.

Independent claim 5, in addition recites the scheduling of tasks to execute in different time partitions. This element is not shown in either reference. Atlas does not contemplate different time partitions. Ezer processes different types of data in the different time partitions, utilizing different system resources to do so. It also lacks any teaching of scheduling tasks to execute in different time partitions. Tasks are preassigned their time partitions based on what type of data they involve. There is no specific teaching of scheduling the tasks.

Independent claim 6 collects unscheduled execution time from one time partition and allocates unscheduled execution time to a task in another time partition. This type of scheduling is clearly not described in the references, and is not addressed in the Final

Office Action. As described above, in Ezer, a task runs in one of the partitions based on the type of data it involves. The time partitions appear fixed, so slack from one partition can not be used in another partition, without changing the timing of the partitions. This is clearly not contemplated by either reference, as Atlas does not deal with time partitions. Thus, it is not possible to combine these references to arrive at the claimed invention, and the rejection should be reversed.

Claim 7 depends from claim 6 and in addition recites that the tasks that are allocated slack are aperiodic, non-essential tasks. Atlas appears to be limited to only a static set of execution threads, i.e. a fixed set of recurring tasks without any new periodic tasks being activated and without any periodic tasks being deactivated. However, claim 2 recites that slack is allocated to aperiodic, non-essential tasks. The Final Office Action asserts that this is taught in Atlas at page 4 lines 19, 23-24, and 36-37. Atlas, references aperiodic tasks in passing comments about other work, but specifically indicates that “We use slack stealing to admit jobs of variable length.” And “For SSJAC, we do not use any aperiodic servers;” in section 3.2, paragraphs 1 and 3 on page 4. Thus, it does not teach the limitations of claim 2 and the rejection should be reversed.

Claim 8 depends from claim 7 and further recites that tasks are allocated slack are from the group consisting of new non-essential tasks and enhancements to essential tasks. The Final Office Action indicates that this is shown in Atlas at page 4 lines 19, 23-24, and 36-37. No such teaching is seen at the referenced lines. Only periodic jobs of variable length appears to be referenced, as well as various priorities of jobs. Since this element appears to be lacking from the references, the rejection should be reversed.

Claim 9 depends from claim 6 and further recites that both timeline slack and reclaimed slack are determined. The Final Office Action references page 4 lines 19, 23-24, and 36-37, and page 5, lines 13-17. These sections of Atlas have been reviewed, and no reference to those types of slack were found.

Claim 10 is a machine-readable medium version of claim 6 and distinguishes from the references for at least the same reasons.

Claim 11 recites a time partitioned system that determines available slack from timeline slack and reclaimed slack, pools it in a common pool, and allocates slack from

the common pool to tasks. This claim is believed to distinguish the references for at least the same reasons as claim 1, and further because the combination of references does not teach or suggest allocating slack in a time partitioned system from a common pool. Specifically, to do so in Ezer would not make sense, given the data based separation of tasks into different fixed partitions. Thus, the references are not properly combinable, and even if combined, do not teach the invention as claimed. The rejection should be reversed.

Claims 12 and 13 depend from claim 11 and distinguish the references for at least the same reasons. Further, they recite features similar to claims 2 and 3, and distinguish the references for the same reasons such claims distinguish the references.

Claim 14 is a machine-readable medium claim similar to claim 11, and distinguishes the references for at least the same reasons.

Claim 15 is similar to claim 11, and further allocates slack to a task in any time partition. As indicated above, such an allocation scheme would not work with Ezer time partitions, since such partitions are specifically described in Ezer as “For a given amount of time synchronized to the video frame, a specified amount of time is partitioned to perform audio tasks; a specified amount of time is partitioned to process video; and the time left over is partitioned to render graphics.” Col. 2, lines 2-7. Thus, it appears that the times for each partition are fixed, and therefore slack from one is useless in another. The rejection should be reversed.

Claims 16 and 17 depend from claim 15, and in addition distinguish the references in a manner similar to claims 2 and 3.

Claim 18 is a machine-readable medium version of claim 15, and distinguishes the references for at least the same reasons.

Claim 19 is a system claim that recites a processor the executes essential and non-essential tasks, each task having a worst case execution time, and an executive that determines available slack and allocates it to tasks in different time partitions. This claim distinguishes the reference for at least the same reasons as claim 1, and further includes the worst case execution times. No reference to such is found in either of the references. Since this element is missing from the references, the rejection should be reversed.

Claims 20-20 depend from claim 19, and distinguish the references for at least the same reasons.

Claim 20 depends from claim 19, and further references determining slack from the group of timeline slack, reclaimed slack and idle time. It distinguishes from the references for the same reasons as claim 4 and claim 19.

Claim 21 depends from claim 20 and further references maintaining a pool of available slack. It distinguishes for at least the same reasons as claim 20, and further no pool of slack from different time partitions is taught anywhere in the references.

Claim 22 is believed to distinguish the references for at least the same reasons as claim 1, and further because the combination of references does not teach or suggest allocating slack in a time partitioned system from a common pool for use by tasks in any time partition. Specifically, to do so in Ezer would not make sense, given the data based separation of tasks into different fixed partitions. Thus, the references are not properly combinable, and even if combined, do not teach the invention as claimed. The rejection should be reversed.

Claims 23 and 24 depend from claim 19 and distinguish the references for at least the same reasons. Further, they recite allocating slack to tasks that are non-essential, new-nonessential tasks, and enhancements to essential tasks. No reference in the text identified in the Final Office Action in Atlas reference essential and non-essential tasks. Thus, a prima facie case of obviousness has not been established, and the rejection should be reversed.

Claim 25 references a module that assigns priority to tasks. The Final Office Action refers to page 5 of Atlas as showing a module that assigns priority levels to different tasks. The language referred to however, relates to job admission control, as stated in the heading for the referenced section. This section describes how jobs are admitted based on priority, and does not describe how tasks are assigned a priority. Since priority may be defined for a tasks by a programmer, there is not teaching or suggestion in Atlas of a module that assigns priority levels to tasks. Thus, a prima facie case of obviousness has not been established, and the rejection should be reversed.

Claim 26 depends from claim 25, and further describes how the first module

determines available slack for tasks at each priority level. The language in Atlas referenced in the Final Office Action does not describe any mechanism for determining slack at each priority level. Rather, Atlas appears to reference available slack, and uses the available slack each time a job is tested to determine whether it can enter the system and run, as described on page 5. Thus, there is no determination of slack at each priority level, only a total of all slack available is maintained. Since this element is lacking from the references, the rejection should be reversed.

Claim 27 also depends from claim 25, and further allocates available slack to tasks in order or priority. As indicated with claim 26, slack is allocated to each job as it is scheduled to run. Available slack is given away “on a FCFS basis.”, as stated in Atlas on page 5. The Atlas algorithm “does not directly schedule based upon task criticality.” However, a resource threshold T can be used to assign different thresholds based on criticality. This appears to be an indirect way of attempting to inject some prioritization into an algorithm that does not inherently take it into account. Thus, it does not meet the language of claim 27, which expressly allocates slack to tasks in order of priority, and the rejection should be reversed.

Claim 28 indicates that the system is a flight control system. The Final Office Action admits that the references do not disclose a flight control system, but indicates that it would be obvious to include a flight control system in the references because it will increase the field of use for such references. This statement is respectfully traversed. An affidavit or reference is requested to show that this would be within the skill in the art at the time the invention was made.

Claim 29 indicates that the system is a real-time control system. The Final Office Action indicates that this shown in Atlas at “page lines 5-8” Applicant has reviewed lines 5-8 of every page of Atlas, and fails to find a reference to a real-time control system. Thus, it is not believed that Atlas describes this claim element, and particularly not in a time-partitioned system.

Claim 30 depends from claim 19, and further references a single set of slack variables and a single slack table. This rejection is respectfully traversed, as the Examiner appears to be combining a description of prior art systems as described in the

related work section starting on page 2 of Atlas, with the system described by Atlas.

There is no suggestion to combine them. Further, Atlas does not describe a single slack table, but merely refers to available slack. There is a mention of potentially storing “level-I inactivity time for every τ_{ij} . If the values are stored in a table, the table must cover a hyperperiod of jobs.....In our implementation, we calculate each job slack when that job is released.” This language does not teach the use of single slack table, and also does not reference a single set of slack variables. As such, the rejection should be reversed.

Serial No.: 09/715,834

Filing Date: December 29, 2000

Title: METHODS AND APPARATUS FOR SHARING SLACK IN A TIME PARTITIONED SYSTEM

8. SUMMARY

Applicant believes the claims are in condition for allowance and requests withdrawal of the rejections to claims 1-30. Reversal of the Examiner's rejections of claims 1-30 in this appeal is respectfully requested.

Respectfully submitted,

PAMELA A. BINNS

By his Representatives,

SCHWEGMAN, LUNDBERG, WOESSNER &
KLUTH, P.A.

P.O. Box 2938

Minneapolis, MN 55402

Date 8-25-2005

By



Bradley A. Forrest


Reg. No. 30,837

CERTIFICATE UNDER 37 CFR 1.8: The undersigned hereby certifies that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail, in an envelope addressed to: Mail Stop Appeal Brief, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on this 25 day of August, 2005.

CANDIS BUENDING

Name

Signature



APPENDIX I

The Claims on Appeal

1. (Original) In a data processing system executing tasks in different time partitions, a method of scheduling tasks comprising:
determining available slack; and
allocating slack to tasks in different time partitions.
2. (Original) The method of claim 1 wherein the tasks that are allocated slack are aperiodic, non-essential tasks.
3. (Original) The method of claim 2 wherein the tasks comprise essential and non-essential tasks, and wherein the tasks that are allocated slack are from the group consisting of new non-essential tasks and enhancements to essential tasks.
4. (Original) The method of claim 1 wherein in determining, both timeline slack and reclaimed slack are determined.
5. (Original) A machine-readable medium having instructions stored thereon capable of causing a processor to carry out a method, the method comprising:
scheduling tasks to execute in different time partitions;
determining available slack; and
allocating slack to tasks in different time partitions.
6. (Original) In a data processing system executing tasks in different time partitions, a method of scheduling tasks comprising:
collecting unscheduled execution time from at least one time partition; and,

allocating the unscheduled execution time to a task in another time partition.

7. (Original) The method of claim 6, wherein the task in the other partition is an aperiodic, non-essential task.

8. (Original) The method of claim 7, wherein the tasks comprise essential and non-essential tasks, and wherein the task in the other partition is from the group consisting of new non-essential tasks and enhancements to essential tasks.

9. (Original) The method of claim 6, wherein in collecting unscheduled execution time, both timeline slack and reclaimed slack are collected.

10. (Original) A machine-readable medium having instructions stored thereon capable of causing a processor to carry out a method, the method comprising:

- scheduling tasks to execute in different time partitions;
- collecting unscheduled execution time from at least one time partition; and
- allocating the unscheduled execution time to a task in another time partition.

11. (Original) In a time-partitioned system executing essential and non-essential tasks, a method of scheduling tasks comprising:

- determining available slack from the group consisting of timeline slack and reclaimed slack;
- pooling available slack in a common slack pool; and
- allocating slack from the common slack pool to tasks.

12. (Original) The method of claim 11, wherein in allocating, slack is allocated to non-essential tasks.

13. (Original) The method of claim 11, wherein in allocating, slack is allocated to a task from the group consisting of new non-essential tasks and enhancements to essential tasks.
14. (Original) A machine-readable medium having instructions stored thereon capable of causing a processor to carry out a method, the method comprising:
 - scheduling tasks to execute in different time partitions;
 - determining available slack from the group consisting of timeline slack and reclaimed slack;
 - pooling available slack in a common slack pool; and
 - allocating slack from the common slack pool to tasks.
15. (Original) In a time-partitioned system executing essential and non-essential tasks, a method of scheduling tasks comprising:
 - determining available timeline slack;
 - determining available reclaimed slack;
 - pooling available timeline and reclaimed slack; and
 - allocating slack to a task in any time partition.
16. (Original) The method of claim 15, wherein in allocating, slack is allocated to a non-essential task.
17. (Original) The method of claim 15, wherein in allocating, slack is allocated to a task from the group consisting of new non-essential tasks and enhancements to essential tasks.
18. (Original) A machine-readable medium having instructions stored thereon capable of causing a processor to carry out a method, the method comprising:
 - scheduling tasks to execute in different time partitions;
 - determining available timeline slack;

determining available reclaimed slack;
pooling available timeline and reclaimed slack; and
allocating slack to a task in any time partition.

19. (Previously Presented) A time-partitioned system comprising:
a processor to execute a plurality of tasks, wherein each task of the plurality of tasks is of a task type selected from the group consisting of essential and non-essential, and wherein each task of the plurality of tasks has associated with it at least one worst case execution time; and
an executive to be in communication with the processor and to control dispatching of tasks on the processor, wherein the executive comprises:
a first module that is to determine available slack; and
a second module that is to allocate available slack to tasks in different time partitions.
20. (Previously Presented) The time-partitioned system of claim 19, wherein the first module is to determine available slack by determining slack from the group consisting of timeline slack, reclaimed slack, and idle time.
21. (Previously Presented) The time-partitioned system of claim 20, wherein the first module is to maintain a pool of available slack.
22. (Previously Presented) The time-partitioned system of claim 20, wherein the first module is to maintain a common pool of available slack that can be used by tasks in any time partition.
23. (Previously Presented) The time-partitioned system of claim 19, wherein the second module is to allocate available slack to tasks that are non-essential.

24. (Original) The time-partitioned system of claim 23, wherein the tasks are from the group consisting of new non-essential tasks and enhancements to essential tasks.
25. (Previously Presented) The time-partitioned system of claim 23, wherein the executive further comprises a third module that is to assign different priority levels to tasks.
26. (Previously Presented) The time-partitioned system of claim 25, wherein the first module is to determine available slack for tasks at each priority level.
27. (Previously Presented) The time-partitioned system of claim 25, wherein the second module is to allocate available slack to tasks in order of priority.
28. (Original) The time-partitioned system of claim 19, wherein the system is a flight control system.
29. (Original) The time-partitioned system of claim 19, wherein the system is a real-time control system.
30. (Original) The time-partitioned system of claim 19, wherein the executive comprises a single set of slack variables and a single slack table.

Serial No.: 09/715,834

Filing Date: December 29, 2000

Title: METHODS AND APPARATUS FOR SHARING SLACK IN A TIME PARTITIONED SYSTEM

APPENDIX II

Related Proceedings

None.

APPENDIX III

Evidence

None.